4

Principles of Code Division Multiple Access

4.1 Introduction

CDMA is a scheme by which multiple users are assigned radio resources using DS-SS techniques. Although all users are transmitting in the same RF band, individual users are separated from each other via the use of orthogonal codes. The North American CDMA standard, or IS-95, specifies that each user conveys baseband information at 9.6 Kbps (Rate Set 1), which is the rate of the vocoder output. The rate of the final spread signal is 1.2288 Mcps, resulting in an RF bandwidth of approximately 1.25 MHz.

There can be many 1.25-MHz signals present in the same RF band. To a large degree, the performance of a CDMA system is interference-limited. This means that the *capacity* and *quality* of the system are limited by the amount of interference power present in the band. Capacity is defined as the total number of simultaneous users the system can support, and quality is defined as the perceived condition of a radio link assigned to a particular user; this perceived link quality is directly related to the probability of bit error, or *bit error rate* (BER). This chapter presents those characteristics of a CDMA system that need to be optimized in order to reduce interference and increase quality.

4.2 Capacity

While there are many models of CDMA capacity in the current literature, we present a description of CDMA system capacity using the amount of user

C1

interference in the band. The actual capacity of a CDMA cell depends on many different factors, such as receiver demodulation, power-control accuracy, and actual interference power introduced by other users in the same cell and in neighboring cells.

In digital communication, we are primarily interested in a link metric called E_b/N_0, or energy per bit per noise power density. Chapter 3 reviews the performance of different digital modulation schemes in terms of probability of bit error as a function of E_b/N_0. This quantity can be related to the conventional *signal-to-noise ratio* (SNR) by recognizing that energy per bit equates to the average modulating signal power allocated to each bit duration; that is,

$$E_b = ST \tag{4.1}$$

where S is the average modulating signal power and T is the time duration of each bit. Notice that (4.1) is consistent with dimensional analysis, which states that energy is equivalent to power multiplied by time. We can further manipulate (4.1) by substituting the bit rate R, which is the inverse of bit duration T:

$$E_b = \frac{S}{R}$$

E_b/N_0 is thus

$$\frac{E_b}{N_0} = \frac{S}{RN_0} \tag{4.2}$$

We further substitute the noise power density N_0, which is the total noise power N divided by the bandwidth W; that is,

$$N_0 = \frac{N}{W} \tag{4.3}$$

Substituting (4.3) into (4.2) yields

$$\frac{E_b}{N_0} = \frac{S}{N}\frac{W}{R} \tag{4.4}$$

Equation (4.4) relates the energy per bit E_b/N_0 to two factors: the signal-to-noise ratio S/N of the link and the ratio of transmitted bandwidth W to bit rate R. The ratio W/R is also known as the processing gain of the system.

Here, we consider the reverse-link capacity since in CDMA this is often the limiting link in terms of capacity. Reverse link is the mobile to base station link. We assume that the system possesses perfect power control, which means that the transmitted powers of all mobile users are actively controlled such that at the base station receiver, the received powers from all mobile users are equal. Based on this assumption, the SNR of one user can be written as

$$\frac{S}{N} = \frac{1}{M-1} \tag{4.5}$$

where M is the total number of users present in the band. This is so because the total interference power in the band is equal to the sum of powers from individual users. Figure 4.1 illustrates the principle behind (4.5). Note that (4.5) also ignores other sources of interference such as thermal noise.

We proceed to substitute (4.5) into (4.4), and the result is

$$\frac{E_b}{N_0} = \frac{1}{(M-1)}\frac{W}{R} \tag{4.6}$$

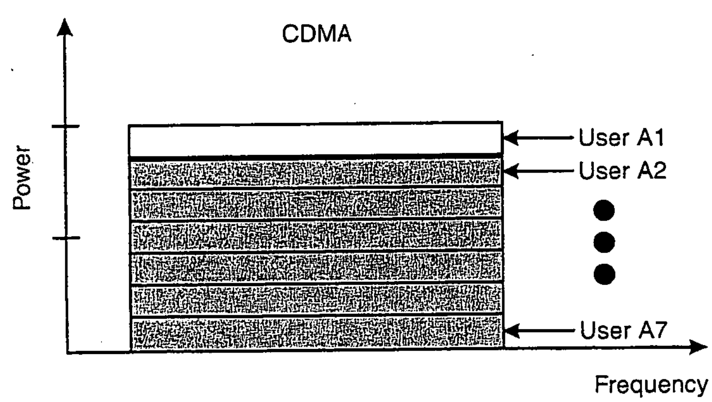

$$\left(\frac{S}{N}\right)_{A1} = \frac{1}{6}$$

Figure 4.1 In CDMA, the total interference power in the band is equal to the sum of powers from individual users. Therefore, if there are seven users occupying the band, and each user is power-controlled to the same power level, then the SNR experienced by any one user is 1/6.

Solving for $(M-1)$ yields

$$M-1=\frac{(W/R)}{(E_b/N_0)} \tag{4.7}$$

Note that if M is large, then

$$M \approx \frac{(W/R)}{(E_b/N_0)} \tag{4.8}$$

4.2.1 Effects of Loading

Equation (4.8) is effectively a model that describes the number of users a single CDMA cell can support. This single cell is omnidirectional and has no neighboring cells, and the users are transmitting 100% of the time. In reality, there are many cells in a CDMA cellular or PCS system. Figure 4.2 shows that a particular cell (cell A) is bordered by other CDMA cells that are supporting other users. Although these other users from other cells are power-controlled by their respective home cells, the signal powers from these other users constitute interference to cell A. Therefore, cell A is said to be *loaded* by users from other cells. Equation (4.6) is modified to account for the effect of loading:

$$\frac{E_b}{N_0}=\frac{1}{(M-1)}\frac{W}{R}\left(\frac{1}{1+\eta}\right) \tag{4.9}$$

where η is the loading factor. η is a factor between 0% and 100%. In the example shown in Figure 4.3, the loading factor is 0.5 resulting in $(1+0.5)$, or a 150% increase of interference above those introduced by home users alone.

The inverse of the factor $(1+\eta)$ is sometimes known as the *frequency reuse factor F;* that is,

$$F=\frac{1}{(1+\eta)} \tag{4.10}$$

Note that the frequency reuse factor is ideally 1 in the single-cell case $(\eta=0)$. In the multicell case, as the loading η increases, the frequency reuse factor correspondingly decreases.

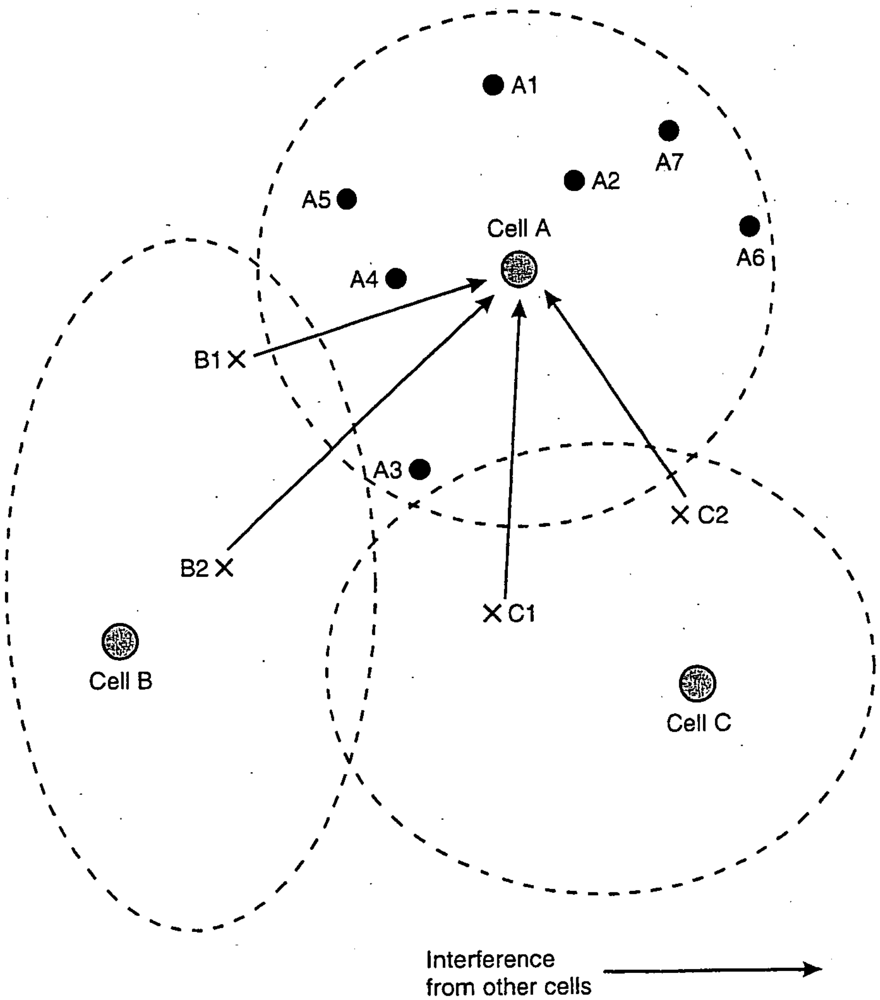

Figure 4.2 Interference introduced by users in the neighboring cell.

4.2.2 Effects of Sectorization

The interference from other users in other cells can be decreased if the cell in question is sectorized. Instead of having an omnidirectional antenna, which has an antenna pattern over 360 degrees, cell A can be sectorized to three sectors so that each sector is only receiving signals over 120 degrees (see Figure 4.4). In effect, a sectorized antenna rejects interference from users that are not within its antenna pattern. This arrangement decreases the effect of loading by a factor of

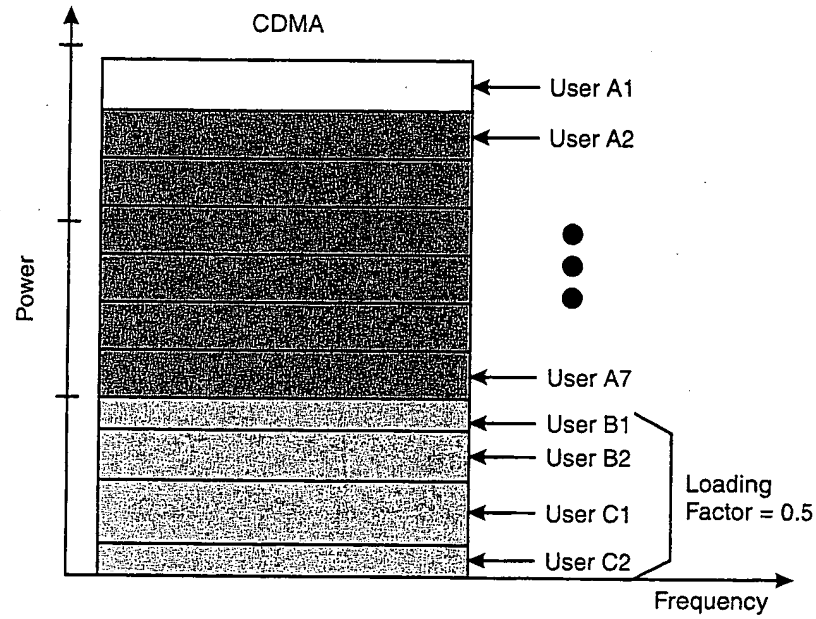

Figure 4.3 Loading factor as perceived by cell A.

approximately 3. If the cell is sectorized to six sectors, then the loading effect is decreased by a factor of approximately 6. This factor is called *sectorization gain* λ.

For one cell, the exact λ is obtained by dividing the total interference power from all directions by the perceived interference powers by the sector antenna; that is,

$$\lambda = \frac{\int_0^{2\pi} I(\theta)d\,\theta}{\int_0^{2\pi}\left(\frac{G(\theta)}{G(0)}\right)I(\theta)d\,\theta} \tag{4.11}$$

where $G(\theta)$ is the horizontal antenna pattern of the sector antenna; $G(0)$ is the peak antenna gain, which is assumed to occur at boresight ($\theta = 0$); and $I(\theta)$ is the received interference power from users of other cells as a function of θ. The integrals in (4.11) are evaluated from 0 to 360 degrees. Equation (4.11) computes the exact sectorization gain, which depends heavily on the antenna gain of the antenna used, as well as on the spatial distribution and distance of

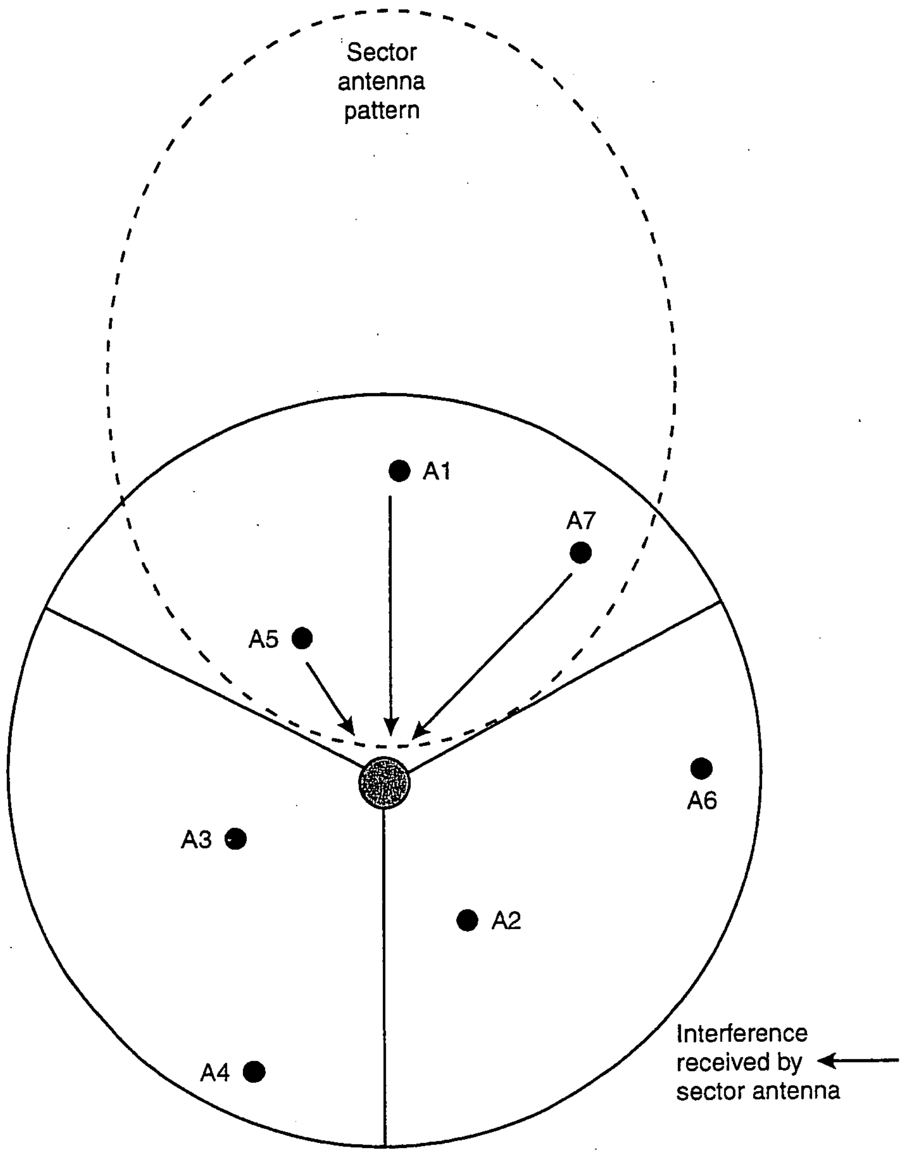

Figure 4.4 A sectorized antenna rejects interference from users not in its antenna pattern.

interfering users in other cells. Note that (4.11) does not take into account the vertical pattern of the sector antenna, the effect of which is quite small in calculating sectorization gain. In reality, λ is typically around 2.5 for three-sector configured systems and 5 for six-sector configured systems.

Equation (4.9) is thus modified to account for the effect of sectorization:

$$\frac{E_b}{N_0} = \frac{1}{(M-1)} \frac{W}{R} \left(\frac{1}{1+\eta} \right) \lambda \tag{4.12a}$$

4.2.3 Effects of Voice Activity

Equation (4.12a) assumes that all users are transmitting 100% of the time. In practice, the vocoder used by the IS-95 system is *variable rate,* which means that the output rate of the vocoder is adjusted according to a user's actual speech pattern. For example, if the user is not speaking during part of the conversation, the output rate of the vocoder is lowered to prevent power from being transmitted unnecessarily. The effect of this variable-rate vocoding is the reduction of overall transmitted power and hence interference. Speech statistics shows that a user in a conversation typically speaks between 40% and 50% of the time. By employing variable-rate vocoding, the system reduces the total interference power by this *voice activity factor.*

Thus, (4.12a) is again modified to account for the effect of voice activity:

$$\frac{E_b}{N_0} = \frac{1}{(M-1)} \frac{W}{R} \left(\frac{1}{1+\eta} \right) \lambda \left(\frac{1}{v} \right) \tag{4.12b}$$

where v is the voice activity factor. Note that the effect of voice activity is to reduce the denominator, or the interference portion of the equation.

Solving (4.12b) for M yields

$$M = 1 + \frac{(W/R)}{(E_b/N_0)} \left(\frac{1}{1+\eta} \right) \lambda \left(\frac{1}{v} \right) \tag{4.13}$$

If M is large, then

$$M \approx \frac{(W/R)}{(E_b/N_0)} \left(\frac{1}{1+\eta} \right) \lambda \left(\frac{1}{v} \right) \tag{4.14}$$

Examining (4.12b), we can draw several conclusions regarding CDMA capacity:

1. Capacity, or number of simultaneous users M, is directly proportional to the processing gain of the system.

2. The link requires a particular E_b/N_0 to attain an acceptable BER and ultimately an acceptable *frame error rate* (FER). Capacity is inversely proportional to the required E_b/N_0 of the link. The lower the required threshold E_b/N_0, the higher the system capacity.
3. Capacity can be increased if one can decrease the amount of loading from users in adjacent cells.
4. Spatial filtering, such as sectorization, increases system capacity. For example, a six-sector cell would have more capacity than a three-sector cell.

4.3 Power Control

4.3.1 Why Power Control?

Power control is essential to the smooth operation of a CDMA system. Because all users share the same RF band through the use of PN codes, each user looks like random noise to other users. The power of each individual user, therefore, must be carefully controlled so that no one user is unnecessarily interfering with others who are sharing the same band.

To illustrate how power control is essential in CDMA, we consider a single cell that has two hypothetical users (see Figure 4.5). We again examine the reverse-link case since this link is often the limiting link in CDMA. User 2 is much closer to the base station than user 1. If there is no power control, both

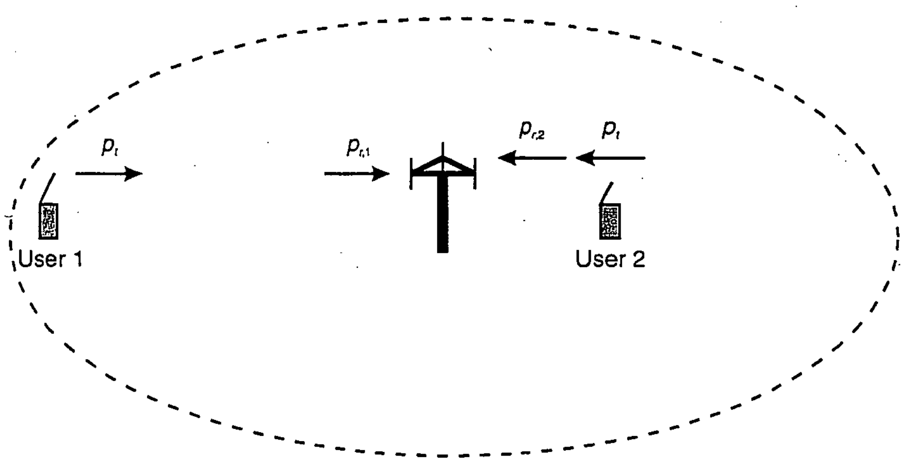

Figure 4.5 A base station with two hypothetical users. Each user is transmitting to the base station a fixed amount of power p_t.

users would transmit a fixed amount of power p_t; however, because of the difference in distance, the received power from user 2, or $p_{r,2}$, would be much larger than the received power from user 1, or $p_{r,1}$. If we assume that the difference in distance is such that $p_{r,2}$ is 10 times more than $p_{r,1}$, then user 1 would be at a great disadvantage.

If the required SNR, $(S/N)_{\text{required}}$, is (1/10), then we can immediately see the disparity between the SNRs of the two users. Figure 4.6 illustrates the point; if we ignore thermal noise, then the SNR of user 2, $(S/N)_2$, would be 10, and the SNR of user 1, $(S/N)_1$, would be (1/10). User 2 has a much higher SNR and thus enjoys great voice quality, but user 1's SNR is barely making the required $(S/N)_{\text{required}}$. This inequity is known as the classic *near-far* problem in a spread-spectrum multiple access system.

The system at this point is said to have reached its capacity. The reason is that if we attempt to add a third user transmitting p_t anywhere in the cell, then the SNR of that third user would not be able to reach the required $(S/N)_{\text{required}}$. Furthermore, if we force a third user onto the system, that third user not only will not attain the required $(S/N)_{\text{required}}$, but also will cause the SNR of user 2 to drop below the required $(S/N)_{\text{required}}$.

Power control is implemented to overcome the near-far problem and to maximize capacity. Power control is where the transmit power from each user is controlled such that the received power of each user at the base station is equal

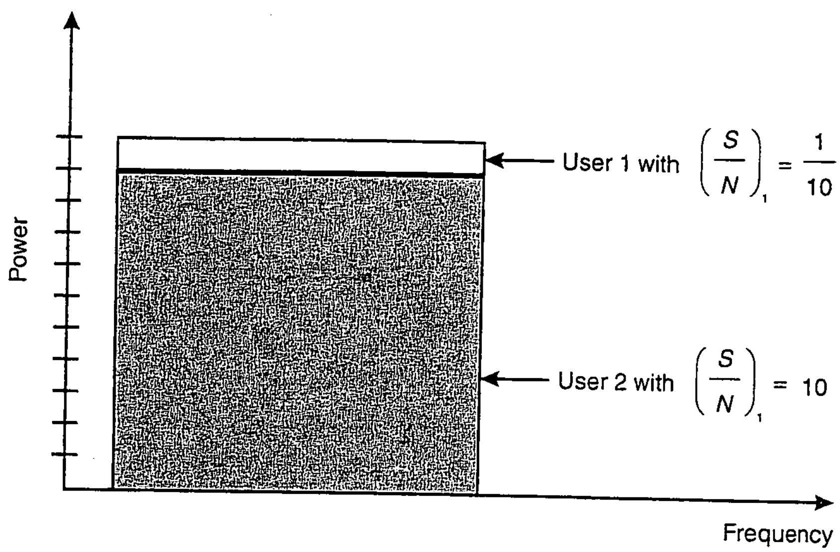

Figure 4.6 Received power from the two users at the base station. User 2 has a much higher SNR than user 1 does.

to one other. Figure 4.7 illustrates the concept. In the cell, if the transmit power of each user is controlled such that the received power of each user at the base station is equal to p_r, then a lot more users can be accommodated by the system. As a continuation of our previous example, if the required SNR $(S/N)_{\text{required}}$ is still (1/10), then a total of 11 users can be supported by the cell. The capacity is maximized with the use of power control (see Figure 4.8).

4.3.2 Reverse Link

4.3.2.1 Access Probes

One problem that has to be immediately solved in power control is the initial mobile transmit power. Before the mobile establishes contact with the base station, the mobile cannot be power-controlled by the base station. Thus, the natural question is when the mobile first attempts to access the base station, what power level should the mobile use to transmit its request? At this point, the base station has not yet made contact with the mobile user, and the base station has no idea as to the location of the mobile user. There are two options: the first option is that the mobile can attempt to access the base station with a high transmit power. Such high power increases the probability that the base station will receive that mobile's access request. However, the disadvantage of a high initial transmit power is that such high power represents interference to

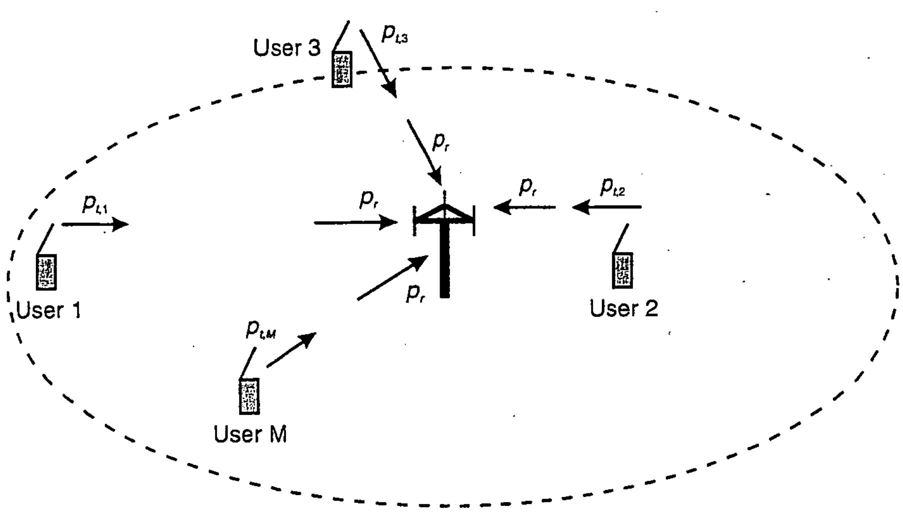

Figure 4.7 With power control, a base station can support many more users. Each user is power-controlled to transmit at different power levels. This is done so that the received powers of individual users are all equal at the base station.

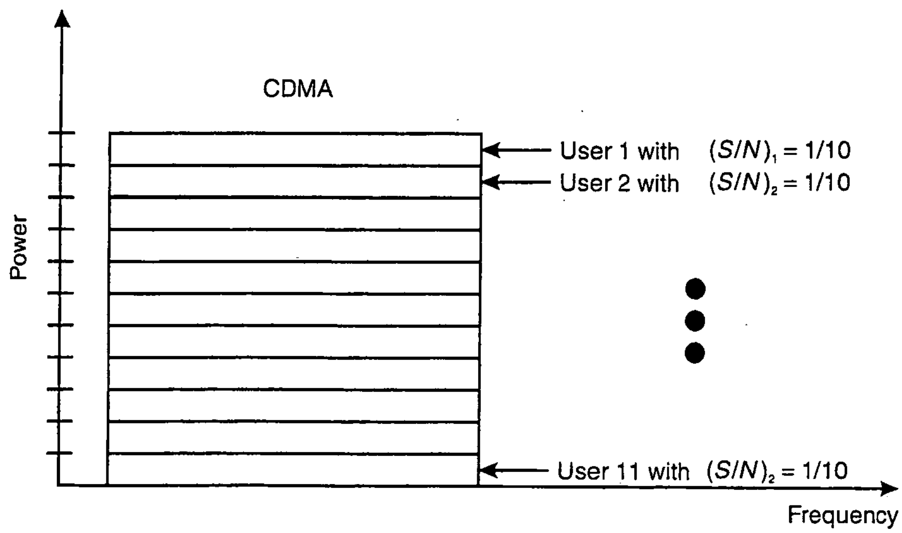

Figure 4.8 Capacity is maximized when the received powers of all users are equal at the base station.

other users currently served by the cell. The second option is that the mobile can request access from the base station with a low transmit power. Such low power decreases the likelihood that the base station will receive the mobile's access request. But the advantage is that this mobile won't cause much interference to other users.

The solution as specified in the IS-95 standard is that when the mobile first attempts to access the system, it transmits a series of *access probes*. Access probes are a series of transmissions of progressively higher power. The mobile transmits its first access probe at a relatively low power, then it waits for a response back from the base station. If after a random time interval the mobile does not receive an acknowledgment from the base station, then the mobile transmits a second access probe at a slightly higher power. The process repeats until the mobile receives an acknowledgment back from the base station. The power difference between the current access probe and the previous access probe is called an *access probe correction* (see Figure 4.9). The step size for a single access probe correction is specified by the system parameter PWR_STEP.

The standard further specifies that the mobile should use the power level it receives from the base station to estimate how much to initially transmit. In other words, if the mobile sees a strong signal from the base station, then it assumes that the base station is nearby and thus transmits initially at a relatively low level. If the mobile sees a weak signal from the base station, then it assumes

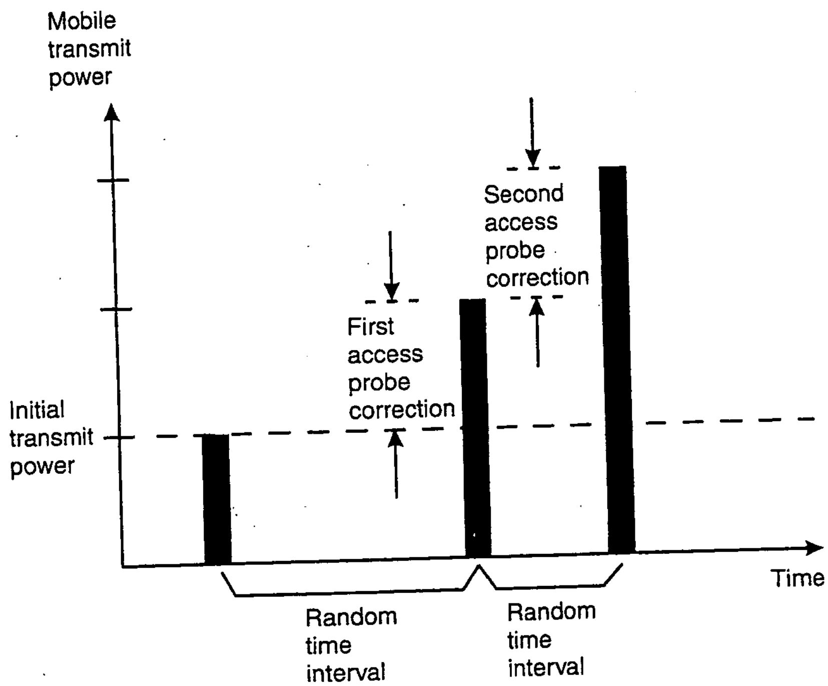

Figure 4.9 A series of access probes by the mobile to access the system. After [1].

that the base station is far away and thus transmits initially at a relatively high level. Knowing the received power from the base station, the mobile can estimate the forward path loss between the base station and itself. If it knows the transmit ERP of the base station, then the mobile would know how much it needs to transmit to compensate for that path loss. In reality, the mobile does not know the actual ERP of the base station, nor does it know how much received power is contributed by other, neighboring base stations. Therefore, a default constant is specified by the standard using generic assumptions of typical loading and base station ERPs. Specifically, the initial transmit power of the mobile, $p_{t,\text{initial}}$ in decibels, should be [1]

$$p_{t,\text{initial}} = -p_r - 73 + \text{NOM\_PWR} + \text{INIT\_PWR} \tag{4.15}$$

As we can see, the default constant is −73 for cellular. A value of −76 is used for PCS systems. The two additional adjustments, NOM_PWR and INIT_PWR, can be set by the system operators for further fine-tuning. The values of these two adjustment factors, NOM_PWR and INIT_PWR, as well

as the parameter PWR_STEP are broadcast by the base station (in the *access parameters message*) and received by the mobile prior to access probe transmission [1]. Upon receiving these two adjustment factors, the mobile uses them in (4.15) to determine its initial transmit power.

4.3.2.2 Open Loop

The process described above is termed *open-loop* power control in that it is purely a mobile-controlled operation and does not involve the base station at all. This open-loop process continues well after the base station has acknowledged the mobile's access request and after the mobile starts to transmit on a traffic channel.

After a call is established, and as the mobile moves around within the cell, the path loss between the mobile and the base station will continue to change. As a result, the received power at the mobile will change and the open-loop power control will continue to monitor the mobile received power p_r and adjust the mobile transmit power according to the following equation [1]:

$$\begin{aligned} p_t = -p_r - 73 + \text{NOM\_PWR} + \text{INIT\_PWR} \\ + \text{ (sum of all access probe corrections)} \end{aligned} \qquad (4.16)$$

where p_t is the continuous open-loop estimate of the mobile transmit power. The difference between (4.16) and (4.15) is that (4.16) contains an additional term specifying the sum of all access probe corrections made during the access probe transmission.

It is important to note that the open-loop power control as specified in (4.16) is based on an estimate of the forward path loss. This power control is used to compensate for slow-varying and log-normal shadowing effects where there is a correlation between the forward-link and reverse-link fades. However, since the forward and reverse links are on different frequencies, the open-loop power control is inadequate and too slow to compensate for fast Rayleigh fading. Note that fast Rayleigh fading is frequency dependent and occurs over every half-wavelength (see Chapter 2 for a discussion of propagation phenomena in a mobile environment). In other words, since fast Rayleigh fading is frequency dependent, we cannot use open-loop power control (which assumes forward path loss is identical to reverse path loss) to compensate for fast Rayleigh fading.

4.3.2.3 Closed Loop

The *closed-loop* power control is used to compensate for power fluctuations due to fast Rayleigh fading. It is closed loop in that the process involves both the

base station and the mobile. Once the mobile gets on a traffic channel and starts to communicate with the base station, the closed-loop power-control process operates along with the open-loop power control. In the closed-loop power control, the base station continuously monitors the reverse link and measures the link quality. If the link quality is getting bad, then the base station will command the mobile, via the forward link, to power up. If the link quality is too good, then there is excess power on the reverse link; in this case, the base station will command the mobile to power down. Ideally, FER is a good indicator of link quality. But because it takes a long time for the base station to accumulate enough bits to calculate FER, E_b/N_0 is used as an indicator of reverse link quality.

The reverse-link closed-loop power control is as follows:

1. The base station continuously monitors E_b/N_0 on the reverse link.
2. If E_b/N_0 is too high (i.e., if it exceeds a certain threshold [2]), then the base station commands the mobile to decrease its transmit power.
3. If E_b/N_0 is too low (i.e., if it drops below a certain threshold [2]), then the base station commands the mobile to increase its transmit power.

The base station sends the power-control commands to the mobile using the forward link. These power-control commands are in the form of *power-control bits* (PCBs). The amount of mobile power increase and power decrease per each PCB is nominally +1 dB and −1 dB.

Because the closed-loop power control is meant to combat fast Rayleigh fading, the mobile's response to these power-control commands must be very fast. For this reason, these PCBs are directly sent over the traffic channel. What actually happens is that bits are *robbed* from the traffic channel in order to send these PCBs. Figure 4.10 shows a simplified block diagram of a portion of the forward traffic channel generation.

The output from the vocoder and input into the convolutional encoder is 9.6 Kbps (at full rate for Rate Set 1). The Rate 1/2 convolutional encoder doubles the baseband rate to 19.2 Kbps. Prior to spreading, the PCBs at 800 bps are multiplexed onto the baseband stream at 19.2 Kbps. The PCBs are integrated into the traffic channel by robbing selected bits from the baseband stream. This way, a separate "channel" at 800 bps (for power-control purposes) exists beneath the traffic channel. The stream of PCBs at 800 bps is therefore called the *power-control subchannel* (PCS). These PCBs are continuously transmitted to the mobile by the base station. Note that since the rate of PCB transmission is 800 bps, a PCB is sent once every (1/800) second, or 1.25 ms

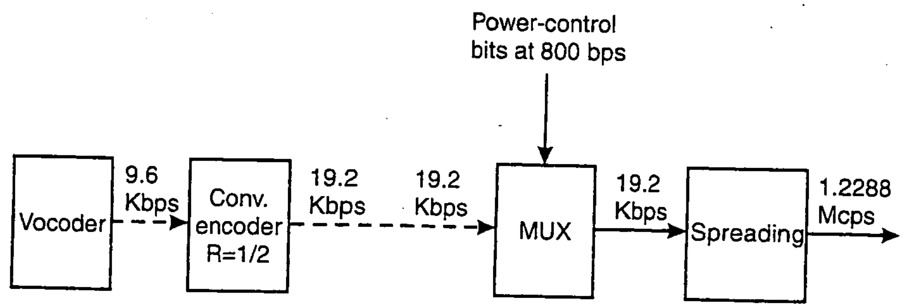

Figure 4.10 In the forward traffic channel, the PCBs at 800 bps are multiplexed directly onto the baseband information stream at 19.2 Kbps.

Both forward-link and reverse-link traffic channel frames are 20 ms in duration. Since one PCB is sent once every 1.25 ms, each traffic channel frame can be divided into (20 ms/1.25 ms) or 16 segments. These segments are called *power-control groups* (PCGs). Since each power-control group is 1.25 ms in duration and the baseband is at a rate of 19.2 Kbps, then each power-control group contains $(19.2 \times 10^3)(1.25 \times 10^{-3}) = 24$ bits. Figure 4.11 illustrates the traffic channel frame structure.

The closed-loop power-control process is illustrated in the example shown in Figure 4.12. For example, for PCG7, the base station measures the SNR or E_b/N_0. The base station compares the measured E_b/N_0 with the threshold. If the measured E_b/N_0 is greater than the threshold, then the base station inserts a PCB of 1 during PCG9 on the forward traffic channel. If the measured E_b/N_0 is less than the threshold, then the base station inserts a PCB of 0 during PCG9 on the forward traffic channel. This process is repeated for every power-control group in the frame.

Since each PCG contains 24 bits (see Figure 4.11), the PCB can be inserted in any one of the first 16 bit positions. The exact location of the PCB in the PCG is determined in a pseudorandom fashion. The PCB bit position is determined by the decimal value of the four most significant bits of the decimator output. The input of the decimator is the long PN code. It is important to recognize that the exact location of the PCB in the PCG is not fixed but pseudorandom.

There are three additional points to mention regarding closed-loop power control.

Power-control bits are not error protected. As we can see from Figure 4.11, the PCBs are multiplexed onto the forward traffic channel *after* the convolutional encoder. Therefore, PCBs are not error protected. This is done to reduce delays that are inherent in decoding and extracting error-protected bits.

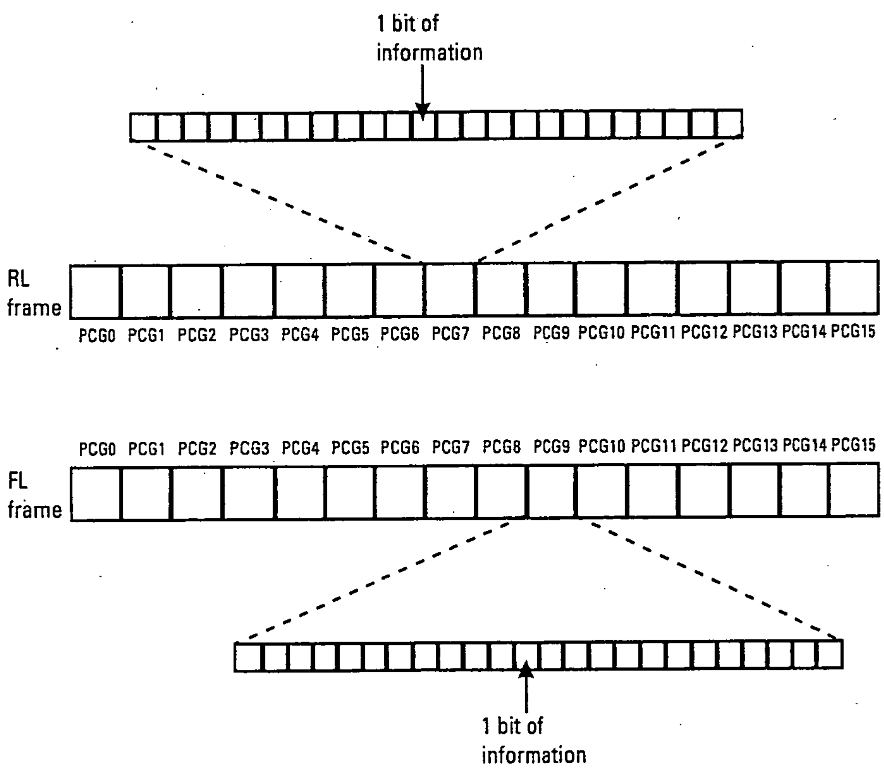

Figure 4.11 The relationship between a traffic channel frame and a PCG.

Remember that the closed-loop power control is used to combat fast Rayleigh fading; the PCBs are not error protected so that the mobile can quickly recover the PCB and adjust its transmit power accordingly. As a result, the probability of bit error for the power-control subchannel may be higher than that of the traffic channel if no special provision is taken.

Closed-loop power control has an inner and an outer loop. We have thus far only described the *inner* loop of the closed-loop power-control process. The premise of the inner loop is that there exists a predetermined SNR threshold by which power-up and power-down decisions are made. Since we are always trying to maintain an acceptable FER, and since in a mobile environment there is no one-to-one relationship between FER and E_b/N_0, the E_b/N_0 threshold has to be dynamically adjusted to maintain an acceptable FER. This adjustment of E_b/N_0 threshold (used by the inner-loop power control) is referred to as the *outer* loop of the closed-loop power control (see Figure 4.13). The outer-loop process is not defined by the IS-95 standard, and each infrastructure

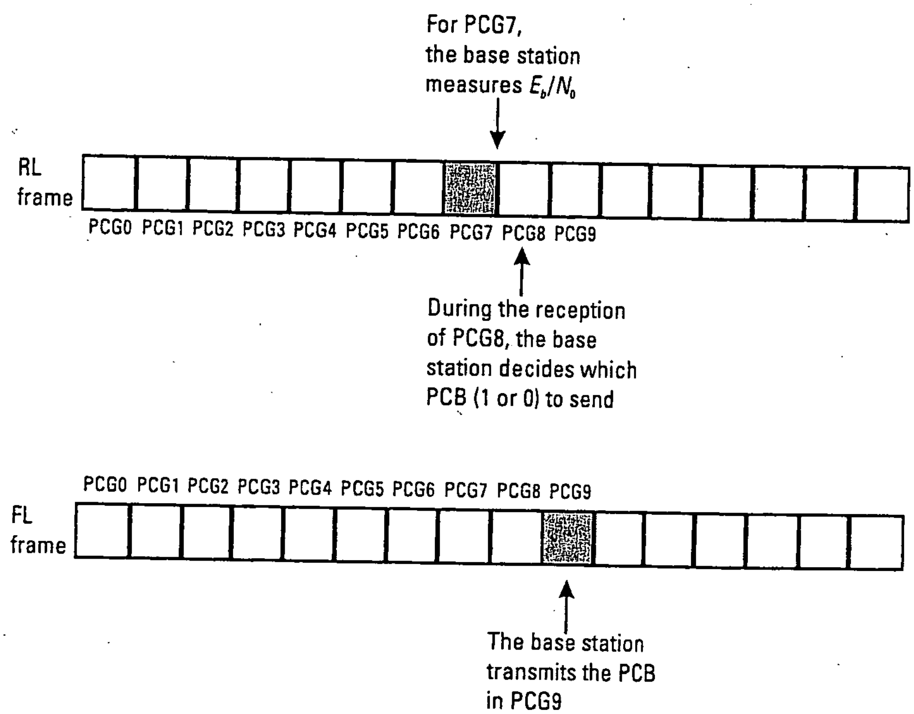

Figure 4.12 Closed-loop power control using PCBs.

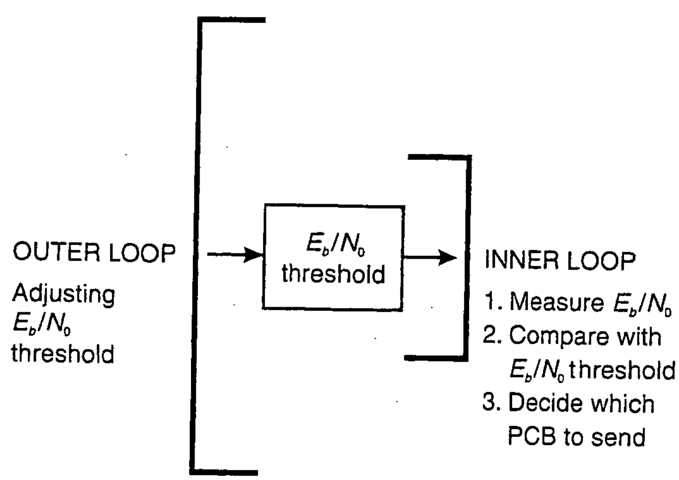

Figure 4.13 Inner and outer loops of the closed-loop power control.

manufacturer is free to implement its own outer-loop algorithms. Note that these algorithms are almost always proprietary to the manufacturer.

The final point concerns *closed-loop power control during soft handoff.* The IS-95 CDMA system utilizes *soft* handoff when a mobile moves between two or three cells. Soft handoff is the process by which a mobile maintains connection with two or three base stations as it transitions between them. During soft handoff, the mobile receives traffic channel frames from two or three base stations, and on these traffic channels there may be conflicting power-control commands (i.e., one base station may be telling the mobile to power up while the other base station may be telling the mobile to power down). In these situations, the mobile follows this rule: if any one base station commands the mobile to power down, the mobile will power down. The mobile will only power up if all of the base stations involved in soft handoff command the mobile to power up.

4.3.2.4 Open-Loop and Closed-Loop Implementation

The mobile transmit power is therefore a function of the open-loop and closed-loop power control of the system. Equation (4.16) can be modified to include the closed-loop power correction; that is,

$$\begin{aligned} \mathrm{p}_t = &-\mathrm{p}_r - 73 + \mathrm{NOM\_PWR} + \mathrm{INIT\_PWR} \\ &+ \text{(sum of all access probe corrections)} \\ &+ \text{(closed - loop correction)} \end{aligned} \tag{4.17}$$

Figures 4.14(a) and 4.14(b) show one implementation of the reverse-link power-control scheme [3]. For the closed-loop power control, the base station has the entire outer loop as well as part of the inner loop; the mobile has the other part of the inner loop. For the open-loop power control, the entire open-loop portion resides in the mobile.

In Figure 4.14(a), the base station receives the reverse-link signal from the mobile. The base station first demodulates the signal and estimates the FER of the reverse link. This information on the reverse-link frame quality is fed into a threshold computer, which adjusts the E_b/N_0 threshold based on the received frame quality. At the same time, the base station also makes an E_b/N_0 estimate of the reverse link. The E_b/N_0 threshold and the E_b/N_0 estimate are then compared. If the estimate is greater than the threshold, then the link E_b/N_0 is higher than what is needed to maintain a good frame quality; a PCB of 1 is thus sent to command the mobile to power down. If the estimate is less than the threshold, then the link E_b/N_0 is lower than what is needed to maintain a good frame quality; a PCB of 0 is thus sent to command the mobile to power up.

The PCBs are multiplexed onto the forward traffic channel and transmitted to the mobile.

On the mobile side (see Figure 4.14(b)), the mobile receives the forward-link signal. It recovers the PCB and, based on the PCB, makes a decision to power up by 1 dB or to power down by 1 dB. The decision is the closed-loop correction. The correction is combined with the open-loop terms, and the combined result is fed to the transmitter so that it can transmit at the proper power level.

4.3.3 Forward Link

Ideally, power control is not needed in the forward link. The reason is that the base station is transmitting all the channels coherently in the same RF band. As Figure 4.15 shows, if thermal noise and background noise are negligible, then all the users fade together as the composite signal arrives at the mobile. However, in real life, one particular mobile may be nearby a significant jammer and experience a large background interference, or a mobile may suffer a large path loss such that the arriving composite signal is on the order of the thermal noise. Thus, forward power control is still needed. In general, however, the power-control requirement for the forward link is not as stringent as that for the reverse link.

The IS-95 standard specifies that the mobile has to report back to the base station the quality of the forward link. The mobile continuously monitors the FER of the forward link, and it reports this FER back to the base station in a message called the *power measurement report message* (PMRM). It may send this report in one of two ways: one is that the mobile periodically reports the PMRM, and the other is that the mobile reports the PMRM only if the FER exceeds a certain threshold. The base station, knowing the quality of the forward link, may then adjust its transmit power to that particular mobile. The exact algorithm of power allocation is again up to the individual infrastructure manufacturer. The process is almost always proprietary to the particular manufacturer.

4.4 Handoff

In a mobile communications environment, as a user moves from the coverage area of one base station to the coverage area of another base station, a handoff must occur to transition the communication link from one base station to the next. The CDMA system as defined by IS-95 supports different handoff processes.

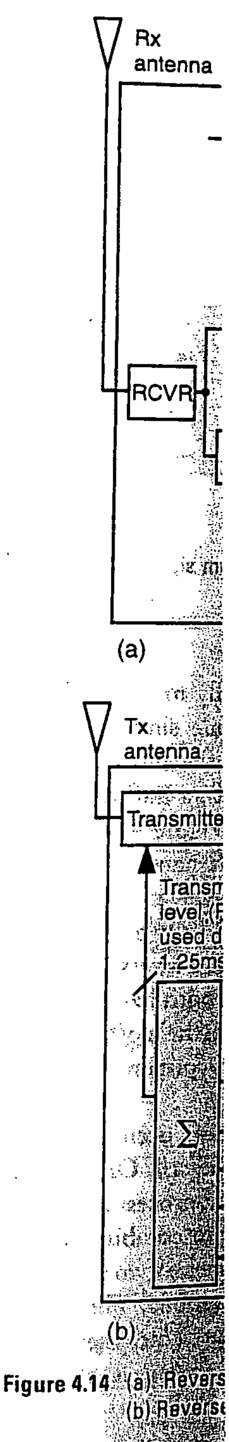

Figure 4.14 (a) Revers
(b) Reverse

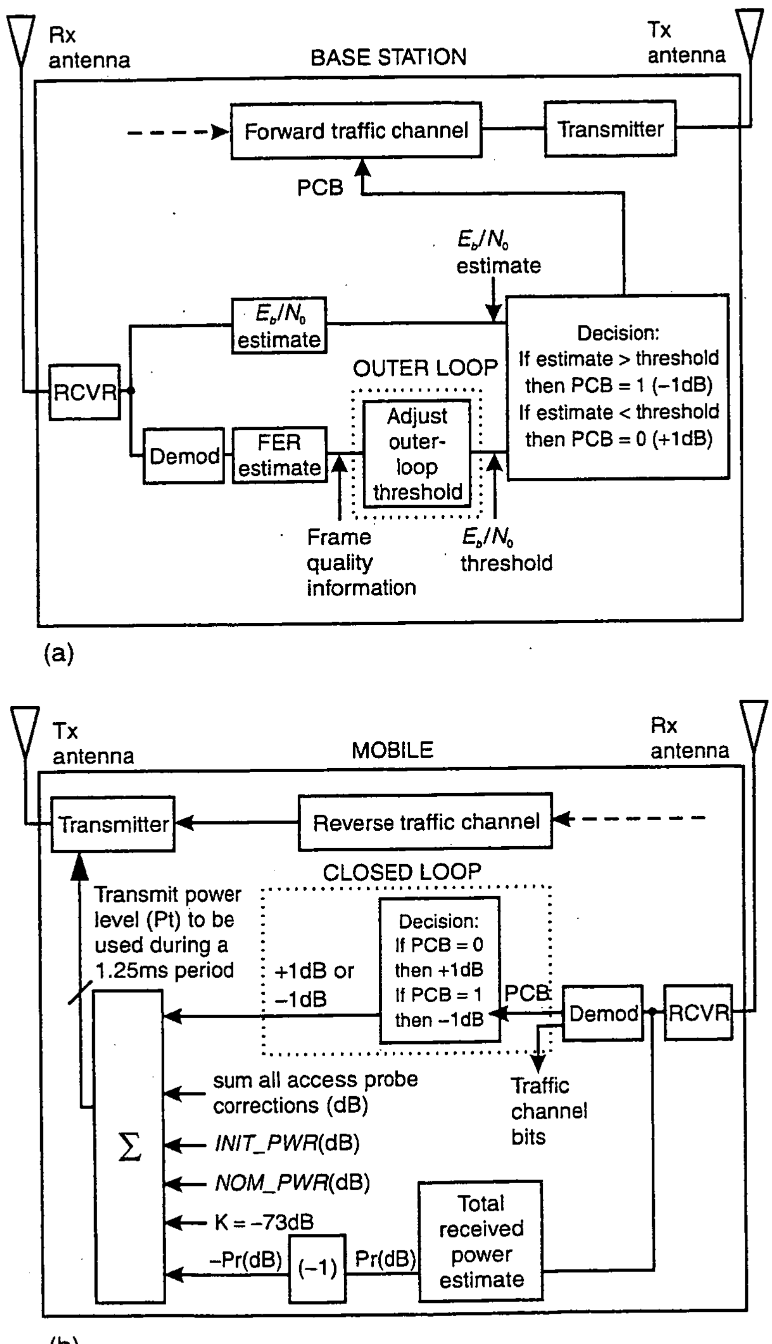

Figure 4.14 (a) Reverse-link power-control functions carried out by the base station; (b) Reverse-link power-control functions carried out by the mobile. After [3].

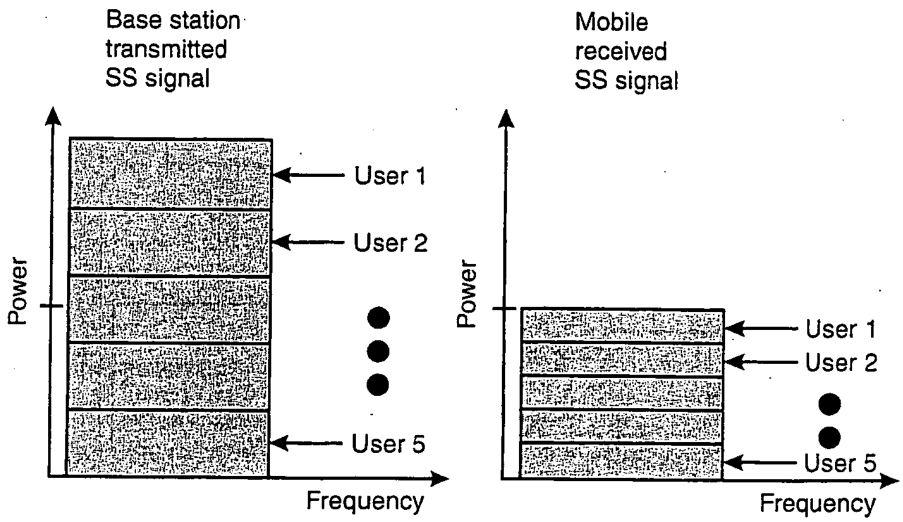

Figure 4.15 All users fade together as the composite spread-spectrum signal travels from the base station to the mobile.

The first is the *soft handoff.* In Section 4.3.2, we briefly mentioned that CDMA uses soft handoff where, during handoff, a mobile simultaneously maintains connection with two or three base stations. As the mobile moves from its current cell (source cell) to the next cell (target cell), a traffic channel connection is simultaneously maintained with both cells. Figure 4.16(a) and Figure 4.16(b) illustrate the simultaneous links during soft handoff. On the forward link (see Figure 4.16(a)), the mobile uses the rake receiver to demodulate two separate signals from two different base stations. The two signals are combined to yield a composite signal of better quality. On the reverse link (see Figure 4.16(b)), the mobile's transmit signal is received by both base stations. The two cells demodulate the signal separately and send the demodulated frames back to the *mobile switching center* (MSC). The MSC contains a selector that selects the best frame out of the two that are sent back.

The second is the *softer handoff.* This type of handoff occurs when a mobile transitions between two different sectors of the same cell. On the forward link, the mobile performs the same kind of combining process as that of soft handoff. In this case, the mobile uses its rake receiver to combine signals received from two different sectors. On the reverse link, however, two sectors of the same cell simultaneously receive two signals from the mobile. The signals are demodulated and combined inside the cell, and only one frame is sent back to the MSC.

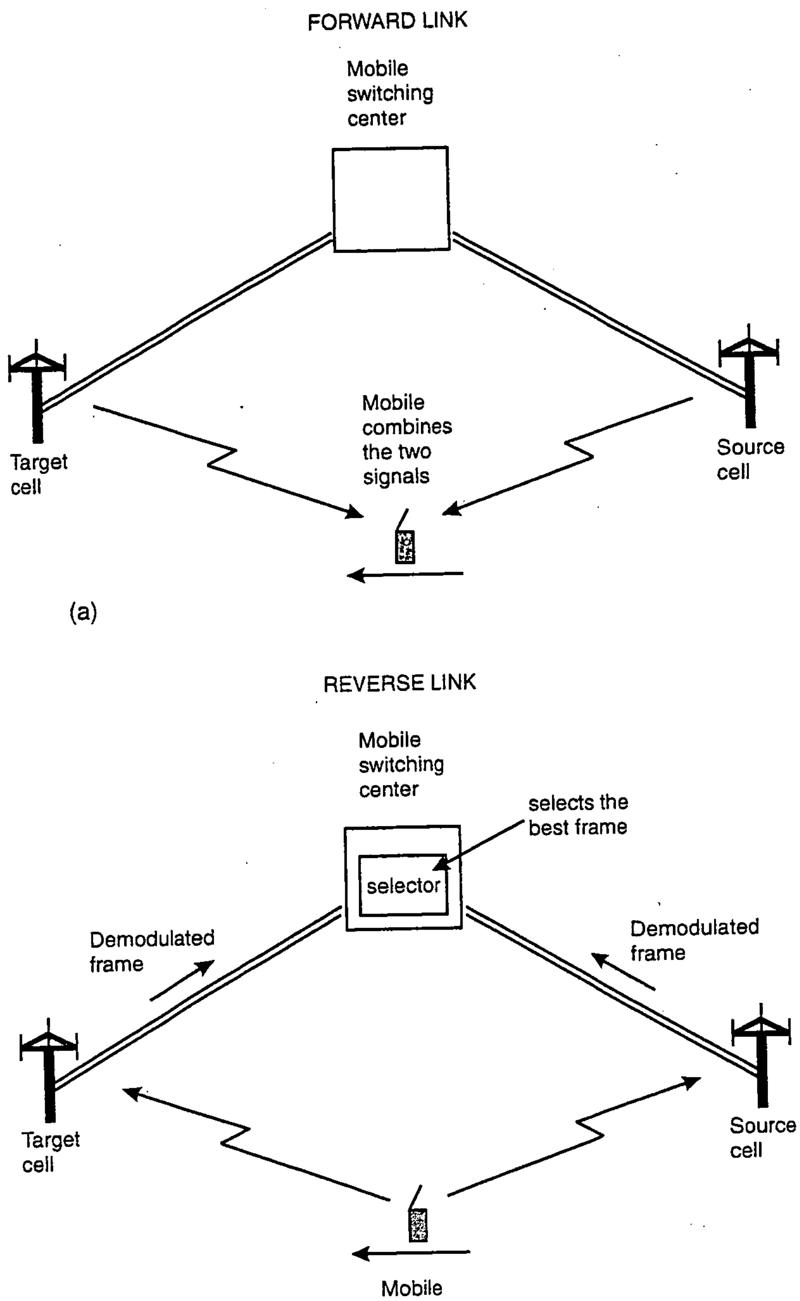

Figure 4.16 (a) Soft handoff between two base stations—forward link; (b) soft handoff between two base stations—reverse link.

Jser 1
Jser 2
User 5
juency
ıvels from

ned that
aneously
le moves
: channel
6(a) and
. On the
demodu-
gnals are
link (see
stations.
ıodulated
a selector

s when a
ı the for-
as that of
ne signals
sectors of
he signals
sent back

The third is the *hard handoff.* The CDMA system uses two types of hard handoffs. *CDMA-to-CDMA handoff* occurs when the mobile is transitioning between two CDMA carriers (i.e., two spread-spectrum channels that are centered at different frequencies). This hard handoff can also occur when the mobile is transitioning between two different operators' systems. CDMA-to-CDMA handoff is sometimes called *D-to-D handoff.* On the other hand, *CDMA-to-analog handoff* occurs when a CDMA call is handed down to an analog network. This can occur when the mobile is traveling into an area where there is analog service but no CDMA service. CDMA-to-analog handoff is sometimes called *D-to-A handoff.*

Before we describe the soft handoff processes in detail, it is important to note that each sector in a CDMA system is distinguished from one another by the pilot channel of that sector. As Figure 4.17 shows, the pilot channel is one of the four logical channels—pilot, paging, sync, and traffic channels—on the forward link. The pilot channel serves as a "beacon" for the sector and aids the mobile in acquiring other logical channels of the same sector. There is no information contained in the pilot other than the short PN code with a specific *offset* assigned to that particular sector. Remember from Chapter 3 that a PN sequence with an offset becomes another PN sequence, and this offset PN sequence is orthogonal to the original sequence. The PN code transmitted on the pilot channel uses this quality to distinguish itself from other sectors and

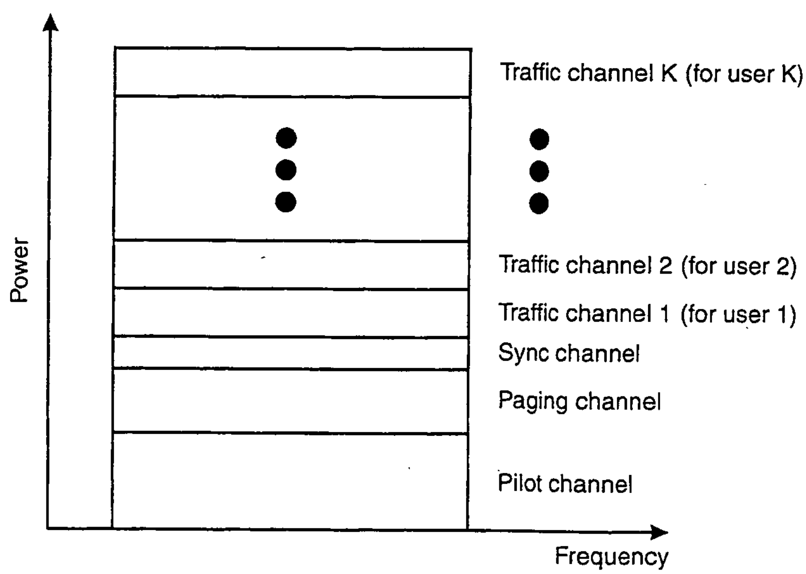

Figure 4.17 CDMA forward link spread spectrum (SS) signal.

other base stations. The offset of a PN sequence associated with a particular sector is specified in the parameter PILOT_PN for that sector.

A special term is used to describe the SNR of the pilot channel: energy per chip per interference density, or E_c/I_0. The energy per chip E_c is different from energy per bit E_b in that "chips" refer to PN sequences that are spread. Since there is no baseband information contained in the pilot channel, the pilot is not despread and bits are not recovered. Therefore, in order to describe the signal strength of the pilot channel, the raw SNR, or E_c/I_0, is used. Note that since the pilot is not despread, E_c/I_0 remains less than 1 most of the time.

4.4.1 Set Maintenance

In CDMA, the mobile is an intimate participant in the soft handoff process. The mobile constantly notifies the base station regarding the local propagation condition; the base station makes use of this information to make handoff decisions. This *mobile-assisted handoff* (MAHO) is evident in that the mobile makes a measurement of forward link E_c/I_0 and reports the measurement result to the base station. Since each base station transmits its own pilot on a different PN offset, the E_c/I_0 of a pilot gives a good indication of whether or not the particular sector should be the serving sector for the mobile.

In managing the handoff process, the mobile maintains in its memory four exclusive lists of base station sectors. The sectors are stored in the form of pilot PN offsets of those sectors. These lists are also called *sets*. The four sets are *active* set, *candidate* set, *neighbor* set, and *remaining* set [2].

The active (A) set contains the pilots of those sectors that are actively communicating with the mobile on traffic channels. If the active set contains only one pilot, then the mobile is not in soft handoff. If the active set contains more than one pilot, then the mobile is maintaining connection with all those sectors on separate traffic channels. The base station ultimately controls the handoff process because a pilot can only be added to the active set if the base station sends a *handoff direction message* to the mobile and the message contains that particular pilot to be added to the active set. The active set can contain at least six pilots.

The candidate set contains those pilots whose E_c/I_0s are sufficient to make them handoff candidates. This means that if the E_c/I_0 of a particular pilot is greater than the *pilot detection threshold* T_ADD, then that pilot will be added to the candidate set. A pilot is removed from this set and placed in the neighbor set if the strength of that pilot drops below the *pilot drop threshold* T_DROP for more than the duration specified by the *handoff drop timer expiration value* T_TDROP. The candidate set can contain at least six pilots.

Note that a pilot can be removed from the active set and placed in the candidate set if the received *handoff direction message* does not include that particular pilot; and if T_TDROP for that pilot has not expired, the pilot is removed from the active set and placed in the candidate set.

The neighbor (N) set contains those pilots that are in the *neighbor list* of the mobile's current serving sector. Initially, the neighbor set contains those pilots that are sent to the mobile in the *neighbor list message* by the serving base station. In order to keep current all the pilots in the neighbor set, the mobile keeps an aging counter for each pilot in this set. The counter is initialized to zero when the pilot is moved from the active or candidate set to the neighbor set. The counter is incremented for each pilot in the neighbor set whenever a *neighbor list update message* is received. The pilot is moved from this set to the remaining set if the counter exceeds NGHBR_MAX_AGE. The neighbor set can contain at least 20 pilots.

Note that a pilot can be removed from the active set and placed in the neighbor set if the received *handoff direction message* does not include that particular pilot; and if T_TDROP for that pilot has expired, the pilot is removed from the active set and placed in the neighbor set.

The remaining (R) set contains all possible pilots in the system for this CDMA carrier frequency, excluding pilots that are in active, candidate, and neighbor sets. The pilot PN offsets in this set are defined by the parameter pilot increment PILOT_INC. For example, if PILOT_INC is 4, then individual sectors in the system can only transmit pilots with offsets of 0, 4, 8, 12, and so forth. PILOT_INC is sent to the mobile in the *neighbor list message* and *neighbor list update message*.

4.4.2 Handoff Process

In the following example, we examine the handoff process from the source cell to the target cell. As Figure 4.18 shows, the mobile is moving from the coverage area of source cell A to the coverage area of target cell B. The following is a sequence of events during this transition:

1. The mobile here is being served by cell A only, and its active set contains only pilot A. The mobile measures pilot B E_c/I_0 and finds it to be greater than T_ADD. The mobile sends a *pilot strength measurement message* and moves pilot B from the neighbor set to the candidate set.
2. The mobile receives a *handoff direction message* from cell A. The message directs the mobile to start communicating on a new traffic

candidate set comparison threshold T_COMP × 0.5 dB, then the mobile sends a pilot strength measurement message.

4.4.3 Pilot Search

In addition to being spread by the Walsh code, the forward link is also spread by a PN sequence (with a specific PILOT_PN offset) at 1.2288 Mcps. This forward-link signal, like any other signal traveling through a mobile environment, can encounter reflections that result in multipaths. As a result, different pilot signals can arrive at the mobile at different times, and a multipath component of one pilot may arrive a few chips later than its direct-path component. Therefore, search windows are provided to search for pilots that are in the active, candidate, neighbor, and remaining windows. Specifically, the parameter SRCH_WIN_A defines the search-window width used to search for pilots in the active and candidate sets, the parameter SRCH_WIN_N defines the search-window width used to search for pilots in the neighbor sets, and the parameter SRCH_WIN_R defines the search-window width used to search for pilots in the remaining sets. These three parameters are sent to the mobile in the *system parameters message* and *handoff direction message*.

The search window for the active and candidate sets is referenced to the earliest arriving multipath component of the pilot. The mobile should center the search window for each pilot in the active and candidate sets around the earliest arriving usable multipath component of the pilot. For example, if SRCH_WIN_A is defined to be 40 chips, then the mobile searches 20 chips around the earliest arriving multipath component of the pilot. For each pilot in the neighbor and remaining sets, the mobile centers the search window for each pilot around the pilot's PN sequence offset using the mobile's timing reference [1].

References

[1] TIA/EIA IS-95A, "Mobile Station-Base Station Compatibility Standard for Dual-Mode Wideband Spread Spectrum Cellular System," Telecommunications Industry Association.

[2] Qualcomm, *The CDMA Network Engineering Handbook, Vol. 1: Concepts in CDMA*, 1993.

[3] Labedz, G., et al., "Predicting Real-World Performance for Key Parameters in a CDMA Cellular System," *Proc. 46th Annual Vehicular Technology Conference*, IEEE, 1996, pp.1472–1476.

Select Bibliography

Feher, K., *Wireless Digital Communications Modulation and Spread Spectrum Applications*, Englewood Cliffs, NJ: Prentice Hall, 1995.

Gilhousen, K. S., et al., "On the Capacity of a Cellular CDMA System," *IEEE Trans. on Vehicular Technology*, Vol. 40, May 1991, pp. 306–307.

Glisic, S., and B. Vucetic, *Spread Spectrum CDMA Systems for Wireless Communications*, Norwood, MA: Artech House, 1997.

Harte, L., *CDMA IS-95 for Cellular and PCS: Technology, Applications and Resource Guide*, New York, NY: McGraw-Hill, 1997.

Kohno, R., R. Meidan, and L. B. Milstein, "Spread Spectrum Access Methods for Wireless Communications," *IEEE Communications Magazine*, Jan. 1995.

Padovani, R. "Reverse Link Performance of IS-95 Based Cellular Systems," *IEEE Personal Communications Magazine*, Third Quarter, 1994.

Peterson, R. L., R. E. Ziemer, and D. E. Borth, *Introduction to Spread-Spectrum Communications*, Englewood Cliffs, NJ: Prentice Hall, 1995.

Pickholtz, R. L., D. L. Schilling, and L. B. Milstein, "Theory of Spread-Spectrum Communications—A Tutorial," *IEEE Trans. on Communications*, Vol. COM-30, No. 5, May 1982, p. 855.

Viterbi, A. J., et al., "Soft Handoff Extends CDMA Cell Coverage and Increases Reverse Link Capacity," *IEEE J. in Selected Areas in Communications*, Vol. 12, No. 8, 1994.

Viterbi, A. J., *CDMA Principles of Spread Spectrum Communication*, New York, NY: Addison-Wesley, 1995.

Viterbi, A. J., A. M. Viterbi, and E. Zehavi, "Performance of Power-Controlled Wideband Terrestrial Digital Communication," *IEEE Trans. on Communications*, Vol. 41, No. 4, 1993.

Viterbi, A. J., and R. Padovani, "Implications of Mobile Cellular CDMA," *IEEE Communications Magazine*, Dec. 1992.